

Assisting the Industrial Internet of Things with Scianta Analytics Extreme Vigilance

Introduction

The web of distributed sensors reporting events and metrics from production scenarios is growing larger every day. This hard-to-manage reality is extremely well-suited to Scianta Analytics' Cognitive Computing approach. Seasonality, long lead times for changes, and irregular physical access make it tough to develop sensible monitoring and alerting systems that are flexible enough to handle the realities of IIoT. Scianta's Extreme Vigilance products are ideally positioned to assist the manufacturing and field deployment support team with solving common business problems.

In many ways, an industrial internet of things deployment could be thought of as just another network of PCs. Systems communicate events and metrics via standard protocols and data is stored for analysis and alerting purposes. However, there are sometimes very significant differences to keep in mind. While an end-user network can be expected to express some of the chaos of human interaction, the Industrial Internet of Things network is often an expression of a designed system. Emergent chaotic behavior may occur in a manufacturing scenario, but it is not

intentional. This principle may become harder to hold as new services are deployed into consumer hands, though. For instance, a network of semi- or fully- autonomous vehicles navigating an urban grid can face chaos on par with that of the Internet's wildest days.



Problem Statements

Alerting as a Trailing Indicator



No one maintaining an IIOT network wants to see chaos! Perhaps there are known key performance and security indicators to watch, or perhaps they can be discovered. Even so, monitoring dozens to hundreds of indicators for a change that might be important is ridiculously expensive. Avoiding problems costs a lot less than fixing them, which is great if you know about problems before they have occurred. Deviation-from-norm alerts are a good starting point, but it takes a much deeper contextual awareness of transactional patterns and seasonality to avoid predictable alert storms.

Where's the Data?



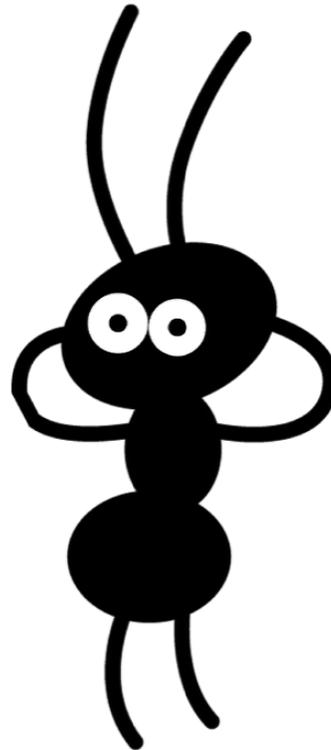
Data acquisition is also a significant challenge in most IIOT environments. While newer equipment might be built to emit signals, a large amount of manufacturing and monitoring equipment was never designed for interaction. Antiquated operating systems, rigid third-party support contracts, and non-negotiable regulatory commitments mean that vulnerabilities and instabilities are built in. Alternate ways can be used to collect data in some organizations, such as passive network capture or access to a database backend managed by another team or vendor; but in other organizations the data scientist can only access extremely limited data sets.

Sound and Fury



Monitoring is hardly new, and in some cases multiple regimes have come and gone; often they're removed because of the

noise generated. While alert noise is arguably better than nothing, a monitoring system that disrupts and exhausts analysts is not helpful. Analysts expect monitoring systems to weigh risk and context appropriately, respond smartly to change, and learn from correction. Above all, the reasons for an alert generation must be clearly discoverable so that logical mistakes can be corrected.



Suggestions

Alerting as a Trailing Indicator



Extreme Vigilance offers multiple ways to detect important problems before an issue. The first to consider is Cognitive Anomaly Detection. The analyst and data scientist work together to define a Data Dictionary which describes the events of interest in terms of actors, assets, and actions. These events are reviewed in a Cognitive Model which automatically determines the band of normalcy for each combination of actors and assets by specific time frames. As new values arrive in the data stream, Extreme Vigilance qualitatively measures the fit with observed data, emitting signals when measurements are approaching or breaching calculated thresholds.

These emitted signals can be used to trigger incident alerts of course, but analysts are also able to define crisp rule sets in the Cognitive Rules Engine to describe known compliance issues and take contextual facts into account. For instance, a rule could be written to state that Inventory must not drop below 100 units. Additionally, rules could be written to express growing concern when inventory is trending downward, or trending downward rapidly, or trending downward very rapidly. Extreme Vigilance calculates reasonable meanings for “very” and “rapidly”, while respecting the crisp limit of 100 units.

Another excellent way to alert ahead of issues is to use Actor and Peer Analyses. These techniques review recent behavior of resources in order to determine how well the behavior matches with past behavior or the behavior of similar resources. If recent measured values are anomalous, signals are

emitted for analysis. Comparing an Actor to itself or to its cohort uncovers subtle variations that may not be visible in threshold anomaly monitoring due to a gradual slope. For instance, peer analysis of smart meters or smart thermostats may be used to uncover insulation problems, while transactional analysis of driving routes could be used to find unmapped road blockages.

Lastly and most powerfully, Extreme Vigilance can model the transactional behavior of an Actor. While metrics and counts may be used to indicate technical problems, alterations in the way that people and systems interact are indications at a business level. Transactional analysis plus anomaly detection opens the door to a better understanding of operations across the board.

Where’s the Data?



Missing data is an insurmountable problem for many data analysis systems, leading to accuracy challenging techniques like interpolation and synthetics. Scianta’s use of the Splunk platform makes it possible to glean high amounts of value from incomplete, indirect, and disparate data streams. Splunk’s rich add-on ecosystem enables access to raw packet capture, infrastructure logs and metrics, database tables, and APIs. While direct access is certainly preferred, Splunk makes it possible to monitor a system indirectly via its impact on infrastructure. In turn, Scianta’s behavior analytics can then operate on these indirect signals.

Sound and Fury



Scianta Extreme Vigilance puts significant effort into producing quality alerts. Signals are generated from a wide variety of rule matches, anomaly detections, and analysis results, but these signals are all weighted. Signal intensity weights are scaled by the severity of the breach, trust level of the model, and criticality of the resources. Analysts can then dial the system's propensity to alert up and down based on their resources and trust in the system's accuracy.

Conclusion

Scianta's Cognitive Computing approach provides transactional analysis and anomaly detection as augmentation to human analysis, purely within the existing Splunk environment. To learn more, see <https://www.scianta.com>

